

## **The Open Science Grid – Support for Multi-Disciplinary Team Science – the Adolescent Years**

**Lothar Bauerdick, Michael Ernst, Dan Fraser, Miron Livny, Ruth Pordes, Chander Sehgal, Frank Würthwein, on behalf of the Open Science Grid**

Fermilab, Brookhaven National Laboratory, Argonne National Laboratory, Computer Sciences Department University of Wisconsin Madison, University of California San Diego.

Email: [ruth@fnal.gov](mailto:ruth@fnal.gov)

**Abstract.** As it enters adolescence the Open Science Grid (OSG) is bringing a maturing fabric of Distributed High Throughput Computing (DHTC) services that supports an expanding HEP community to an increasingly diverse spectrum of domain scientists. Working closely with researchers on campuses throughout the US and in collaboration with national cyberinfrastructure initiatives, we transform their computing environment through new concepts, advanced tools and deep experience. We discuss examples of these including: the pilot-job overlay concepts and technologies now in use throughout OSG and delivering 1.4 Million CPU hours/day; the role of campus infrastructures- built out from concepts of sharing across multiple local faculty clusters (made good use of already by many of the HEP Tier-2 sites in the US); the work towards the use of clouds and access to high throughput parallel (multi-core and GPU) compute resources; and the progress we are making towards meeting the data management and access needs of non-HEP communities with general tools derived from the experience of the parochial tools in HEP (integration of Globus Online, prototyping with IRODS, investigations into Wide Area Lustre). We will also review our activities and experiences as HTC Service Provider to the recently awarded NSF XD XSEDE project, the evolution of the US NSF TeraGrid project, and how we are extending the reach of HTC through this activity to the increasingly broad national cyberinfrastructure. We believe that a coordinated view of the HPC and HTC resources in the US will further expand their impact on scientific discovery.

### **1. Where We Are Today**

The Open Science Grid has completed the its 2006-2012 project cycle. The Department Of Energy (DOE) and the National Science Foundation (NSF) have recently awarded the project funds for the next five years, 2012-2017, with reduced budget. While planning for this next phase, the OSG leadership continues to ensure sustained production services, retention of the expert and excellent staff, and smooth transitioning of the organization and scope.

The DOE and NSF awards have endorsed the OSG leadership and organization. For us, this new phase is likened to “adolescence”. We have survived and grown through our baby-steps (Trillium and Grid2003) and our childhood (OSG). During the next five years we aim to mature and consolidate our DHTC environment and processes such that in 2017 the OSG virtual facility can be more self-sustaining in continuing its production services to science and more agile and energetic in its delivery of new capabilities, scales, and user effectiveness in extreme- and exascale- collaborative research.

In this overview paper we include give two key areas of activity today, the evolution of the organization, the analysis of accomplishments[1] to the goals in the 2006 Consortium charter, and discuss several focuses of the new project. Other CHEP 2012 papers describe more specific aspects of the software, technologies, operations and use of the OSG fabric of services. Papers from previous CHEPs describe the services offered by the OSG.

#### ***1.1. Effective Services: e.g. A Key Transformation - Job Overlays***

Over the last few years the “job overlay” principles and technologies have expanded their reach and role for nearly all communities depending on the OSG. Over fifty-percent of all work on the DHTC

fabric of services is done through the GlideinWMS system[2], an overlay on the CONDOR[3] system. OSG supports submission “factories” at UCSD and Indiana in support of all users, and at CERN for global CMS. Submissions through the ATLAS job overlay system, PANDA[4], account for another twenty-five-percent of the usage.

For the users, job overlays have given a tremendous boost in usability of the services, with easier access to cycles across multiple resources, end-user errors significantly reduced and configuration opportunities which allow for more effective provisioning of the available resources. For the OSG, there has been an increase in effectiveness with the response to site failures being handled centrally by experts, faster fixing of issues and deployment of new capabilities, and increased ability to configure and tune of the job submission engines through the Factory “software as a service”. Further advantages have been the extensions to well defined mechanisms in support provisioning of and access to new types of resources such as commercial clouds.

### ***1.2. An Anchor for Distributed High Throughput Computing: e.g. A Home for Software***

As an anchor for DHTC the OSG offers both consulting and software services in addition to those needed to sustain production computing. One of the commitments is to support software needed by our stakeholders in the event that the software provider is unable to continue development and support. The process has been well exercised for the Bestman[5] storage management software, which was unable to garner sustained funding after 2011. We followed our usual practices for software inclusion and support into the OSG software stack:

- Gather the community needs.
- Assess the impact of supporting the software.
- Identify the owner within OSG and define the Service Level Agreement for the software support.
- Reprioritize OSG efforts to provide the additional services.
- Deploy the operations services and processes in support of the software.
- Transition the responsibility of support from the previous supplier.

Currently we are transitioning the PKI services and technologies needed by our users to OSG from the ESNET DOEGrids CA (at the request of DOE). We are implementing services that depend on a commercial backend from DigiCert. We are also investigating the impact of fall-off of funding support for the core Globus toolkit[6] from which OSG stakeholders depend on GridFTP, GRAM, GSI, and JGlobus)

## **2. Evolving the Organization**

The organization of OSG provides a framework both for delivering agreed upon operational objectives and for allowing community driven extensions in the priorities and deliverables. The OSG eco-system consists of the Consortium[7], the Council, the Project and the Satellites. The Consortium includes all (virtual and physical) organizations that are registered as members – including resources owners, software providers, users etc. The Council represents the Consortium as the governing body of OSG. The Council currently has 20 members (the by-laws specify a maximum of 30), with the process to join requiring sponsorship by an existing member followed by a Council decision (in discussion with the applicant) on the mutual benefits, enthusiasm and commitment of the new participant to be an active and significant contributor to the Consortium vision.

The project is run by the Executive Team, the on the ground active leadership of the OSG, governed by the Council. Within the Executive Team the roles are defined to promote an effective organization for making decisions, prioritizing and managing the program of work. The PI has overall fiscal and program responsibility, the Executive Director is responsible for the deliverables and day-by-day program of work, the Resources Manager has oversight on all resource decisions, the Technical Director leads the distributed high throughput computing conceptual framework and techniques, while the Production Manager oversees key components of the program ensuring effective use of the infrastructure and leading the thrust into the campuses. The Project Manager supports the

Executive Team in tracking and reporting on the program of work, agreements and budgetary matters and the Applications Coordinators provide active bi-directional communication with major stakeholders. We have strengthened the role and responsibilities of the Technology Investigation area to provide engineering contributions as part of the WLCG Technology Evaluation Groups, internally within the OSG, as well as with our federated and satellite partnerships. There has been a transition in the Executive Director and Council Chairs. The leadership is augmented by five Area Coordinators who gather requirements for, define, lead and execute the technical program of work in close cooperation with and overseen by the Executive Team. The Area Coordinators increasing technical and leadership qualities are a significant source of strength of the OSG as it moves through adolescence towards its mature sustaining phase. Mechanisms for changing and/or augmenting the team have been exercised over the past five years, resulting in an increasingly strong and expert set of leaders and succession possibilities. The management plan will be modified to reflect the evolving goals of the changing organization and the thrusts of the overall program.

### **3. Objectives for OSG 2012-2017**

The overarching objectives for the next five years:

- Empower researchers on campus in the use of DHTC. This includes routine and transparent sharing and access to small to large-scale computing resources locally and remotely, using native identities, with full access to data wherever workflows are executed.
- Continue support for ATLAS, CMS, ALICE distributed computing in the US and contributions on behalf of US LHC to WLCG.
- Integrate transparent access to any resources type accessible to OSG users, including GPUs, multi-core, Amazon clouds, and high performance computing centers.
- Increase the number of active OSG communities.
- Effective contributions as an XSEDE Service Provider.
- Contribute to the NSF XD, CIF21 and S2I2 programs, including to the institute program developing centers of excellence on particular cyberinfrastructure topics across all science domains.
- Contribute as part of the evolution of DOE extreme-scale collaborative science computing programs.

Overall we are looking forward to placing a stronger emphasis on support for intra- and inter-campus infrastructures, give ubiquitous access and management of any available and useful resource type, make the necessary extensions to meet the experiment goals of smooth operations after the upcoming LHC shutdown, and extending our active communities and partnerships both at home and abroad.

We give below more information about three key thrusts. Other activities will address the needs of federated identities, ongoing provision of software and production services, operational security, evolution of distributed high throughput computing concepts, and user support.

#### **3.1. The Campuses**

Our ultimate goal is to enable researchers and administrators across all campuses in the US to share resources locally, and to easily and automatically manage access to additional resources remotely on other campuses, facilities and/or clouds, as they need them.

The DHTC principles—diverse resources, dependability, autonomy and mutual trust—that OSG advances and implements at a national level, map well to a campus environment. The goal of our Campus DHTC infrastructure activity is to translate this natural fit into a wide local deployment of high throughput computing capabilities at the nation's campuses, bringing locally operated DHTC services to production in support of faculty and students as well as enabling integration with the evolving national CI[8].

Intra-campus and inter-campus sharing of computing resources enhances scientific competitiveness and, when interfaced to the OSG and XSEDE infrastructures, increases the national computational throughput. To accomplish this we have defined the following comprehensive approach that aims to eliminate key barriers to the adoption of HTC technologies by small research groups on our campuses.

- Support for local campus identity management services removes the need for the researchers to fetch and maintain additional security credentials such as grid certificates.
- An integrated software package that moves beyond current cookbook models that require campus IT teams to download and integrate multiple software components. This package does not require root privileges and thus can be easily installed by a campus researcher.
- Coordinated education, training and documentation activities and materials that cover the potential, best practices and technical details of DHTC technologies.
- A campus job submission point capable of routing jobs to multiple heterogeneous batch scheduling systems, with the initially supported systems being based on Condor flocking.
- An OSG submission point that can route jobs directly to on-campus resources. Hence providing a natural mechanism for existing OSG sites to expand into a campus DHTC infrastructure.

The job submission interfaces presented to campus faculty and students are identical to those already in use today on the OSG infrastructure. Seamless access is starting to be provided to export work and data to OSG, XSEDE, Amazon, and GPUs.

We are currently working with early releases of the software and processes that include enabling the formation of local HTC partnerships and dynamic access to shared local (intra-campus) and remote (inter-campus) resources using campus identities. Working closely with our software providers we integrate a complete set of campus services (i.e.. job submission; accounting; discovery; scheduling; pilot factories) that are easy to install and operate by a small research group. We are also working through the XSEDE Campus programs (campus champions, campus bridging) to understand how best to leverage the total effort.

### ***3.2. Data – from Small Data to Big Data***

Usable, effective management of data remains one of our main challenges in the shared resource environment. While large communities such as the LHC have their own implementations and services, small communities do not have the resources for such luxuries. We are paying attention to several external software developments and working actively within OSG to understand the available use cases, integration, deployment, and support impacts of several emerging technologies. This includes the wide deployment of XROOTD remote I/O services provided by the “Anydata, Anytime, Anywhere” project, Globus Online for the managed transport of files directly or between third parties, wide-area LUSTRE being brought to the table by the ExtENCI satellite, and the IRODS technologies as a means to manage the opportunistically accessible storage at computing sites. Several other CHEP papers describe these technologies in detail. For the purposes of this overview paper we emphasize the importance of these services, our failure to date to meet the needs of the communities, and the intent for some transformation in the next few years in the usability and effectiveness of the solutions.

### ***3.3. Full Support for Additional Resource Types***

Several communities need access to high-core count nodes, clouds, GPUs and other additional resource types. Over the past few years OSG itself and/or collaborating satellites have demonstrated the feasibility of access to and use of such resources by our communities. In the next few years the technologies and services will be moved to regular production support. This transition requires us to ensure that for each resource type the job submission and execution interfaces, administration and operations services, error reporting and fault diagnosis, as well as resource management and policies, are fully integrated into the infrastructure and can be supported by the OSG staff. In some cases the mode of operation is significantly different than the traditional “grid cluster” interfaces we support today, e.g. the identity management, cost model, policy and provisioning for the Amazon EC2 cloud.

We will extend support for provisioning and use of the shared network resources. We will extend our operations services to add community-developed services to monitor the end-to-end network paths to the sites and alert when problems are seen. As the mechanisms to manage the network from the middleware and applications layers improves we will make sure that the data and job services can effectively take advantage of these new capabilities. Finally, the community's increasing data scales and latency reductions needed from shared use of this resource must increasingly be prioritized and policy agreements applied. We will provide mechanisms to configure and tune the overall infrastructure and aggregate throughputs in response to the immediate resource provisioning needs.

#### 4. Strategic Planning and Assessment

The Consortium Charter has recently been revised, after six years, to reflect the changing membership of the OSG and eco-system of DHTC. The changes are in part to address newly identified needs of campus based collaborative computing, the extreme-scale science community and evolution in the capabilities and opportunities in distributed high throughput computing.

Based on the original charter the Council has reviewed our high level accomplishments and lessons learned "report card". These are providing additional input to the strategic and tactical planning for the next five years. We give details below:

	<i><b>Objective</b></i>	<i><b>Accomplishment</b></i>	<i><b>Lack</b></i>	<i><b>Lessons Learned</b></i>
1.	Build a persistent national grid infrastructure for large scale US science the Open Science Grid.	The OSG based infrastructure has been operating continuously from November 2003 when it was instantiated for SC2003 as the "Grid3 infrastructure".	Sustained funding for the infrastructure has not been attained.	Engaging researchers in campus-based contexts was not recognized earlier enough in the project.
2.	Provide a set of goals and an overall infrastructure within which the Grid resources of the different members can be operated coherently and compatibly.	OSG resources have been used by more than 30 different scientific and research groups during the period 2006-2011. There have been two full downtimes of the infrastructure during that time both due to certificate authority issues.	The threshold to entry is regarded by small groups of researchers as too high.	Spend more time understanding the different community needs and cultures and adapt developments to better match the needs.
3.	Computational and application scientists, working together to provide and support the set of facilities, services and infrastructure needed.	OSG management and execution teams are all multi-disciplinary. The OSG Executive Team includes 3 engineers, 1 computer scientist, and 3 application scientists. The 5 areas coordinators are 3 computer scientists and 3 engineers.	The balance of computer scientists' use of and contributions could be improved.	Within the OSG management team we learned how to leverage and coordinate with facility managers of the participating VOs who own the computing and storage resources as well as the systems administration effort.
4.	A structure of management and coordination bodies will oversee and coordinate the work of the Open Science Grid Consortium.	The OSG Consortium developed and agreed upon by-laws, agreed upon a project management plan - that have been adapted as needed.	The co-engagement of the Council members has not been as energetic as hoped for.	
5.	The Open Science Grid will be open to all sciences that have a need for distributed large scale computing and data management, and can bring resources to be federated.	OSG services have benefited researchers from more than 10 science domains. New entrants to the Council include the Structural Biology Grid consortium and the University of Nebraska	The contributions of non-physics communities and organizations was hoped to be.	As noted above, the importance of using the campus context to connect and coordinate with groups whose computing environments are evolving was not recognized until late in the project.
6.	An engineered production quality grid infrastructure will be built and operated	OSG maintains a stable, increasing scale, production infrastructure across more than	Some of the Blueprint has been inconsistent and	Our approach could be improved by more frequently capturing changes in the reference

	<b>Objective</b>	<b>Accomplishment</b>	<b>Lack</b>	<b>Lessons Learned</b>
	in the US.	100 sites. The availability of the core operations services is greater than 99% in most cases. The usage of the infrastructure has steadily increased. The Blueprint established in 2005 has been actively used.	out of date. The updating has been uneven.	architecture and Blueprint. This would also give Satellite development projects a more complete target environment.
7.	Extended internationally through participation in the global grid infrastructure for science.	OSG is recognized as an important contributor to the Worldwide LHC Computing Grid. OSG has also federated partnerships in Colombia and Brazil. Notably Colombia has created an independent NGI modelled after OSG's architecture and using OSG's software stack.		Good documentation and training materials, as well as in person visits and education might reduce effort overall. We developed a framework for re-use of production documentation as well as materials contributed by expert volunteer 'faculty' at OSG site administrator workshops.
8.	The grid will support managed access to large computing and storage resources (up to tens of thousands of CPUs and 100s of petabytes).	OSG today provides access to more than 50,000 job slots over more than 70,000 cores and over 30 petabytes of useable storage through a variety of backend technologies and protocols. It remains a powerful organizational framework that connects thousands of users to these resources on a continuous basis. Much more detail is given in the OSG Annual Reports.	While the OSG provides access to tens of thousands of cores, 100-petabyte scales is yet to be reached in part due to variability in the data distribution and caching strategies of the participating VOs.	We need to pay more attention to showing the value of sharing and fast turnaround to initial results. For small VOs to share they must support staff and infrastructure to answer and address problems from a diverse community with which they are not necessarily familiar. Funding such effort to improve the effectiveness of the whole could be considered.
9.	Create opportunities for educators and students to participate in building and exploiting this grid infrastructure and opportunities for developing and training a scientific and technical workforce.	OSG has held three summer schools. More than 75 students have been educated in distributed high throughput computing concepts and use. Additionally through site administrator workshops more than 70 systems administrators have been trained in best practices for grid service deployment and operations,	The program remained challenged until the job "pilot/overlay" became available for large-scale job submission.	Education, training and outreach activities should in general be closely coupled to the core activities of a project; therefore as the educational mission included students and distributed computing and new science domains, Having robust, common user frameworks with a quick application integration capabilities for new science domains is vital.
10.	Ensure that the U.S. plays a leading role in defining and operating the global grid infrastructure needed for large-scale collaborative and international scientific research.	OSG is recognized as an important contributor to the Worldwide LHC Computing Grid, supports international collaborations and has federated infrastructure partners in Colombia and Brazil.		

## 5. Preparing for Post-Adolescence

As a center for the concepts, practice and evolution of distributed high throughput computing, the OSG depends on computer science research to advance the methods, techniques and to drive the effective technologies to meet the expanding needs of the science community. As input to this we have recently developed a list of computer science research on which the continued health of the OSG services will depend. We give a summary below:

- Transparent usage and ad-hoc/on-demand scheduling of heterogeneous resources, including sensors, mobile devices, laptops, loosely-coupled clusters, clouds, high-end computing systems etc.
- Life-time and end-to-end data management and access across globally distributed and heterogeneous storage resources, including content-managed, global distribution and low-latency transparent access to vast multi-source data sets as well as efficient management of and access to large aggregations of small, diverse, dynamic, data-types.
- Policy based management of the ensemble of resources together with co-scheduling of multiple resource types in complex topologies and coordinating the on-demand provisioning of network resource with data access systems.
- Maturing the robustness of solutions together with definition and implementation of quality and assessment methods.
- Adaptable and flexible use of dynamic many-core computing hardware, GPUs and further new hardware developments.
- Distributed identity and access control frameworks for sharing and use of information and resources including the management and application of dynamic, short- to long-lived collaboration based access control.
- User level usability, flexibility and integration including provisioning of dynamic, integrated end-to-end environments for shared- use and high level, tailored workflow definitions and engines.

## 6. Summary

The participants in the OSG life remain convinced of the value of the conceptual and implementation eco-system of DHTC. We remain committed to bring the benefits of DHTC to the broad community of computationally based research and science. The goals for our virtual facility and our fabric of services are designed to meet the needs of 21<sup>st</sup> century collaborative science. We will continue to track, assess and learn from our ongoing activities to be able to graduate at the end of “high school”. We are happy and excited to continue this unique collaboration and endeavor.

## References

- [1] OSG accomplishments <http://osg-docdb.opensciencegrid.org/cgi-bin/RetrieveFile?docid=1086&extension=pdf>
- [2] Sfiligoi, I., et. al., Making science in the Grid world: using glideins to maximize scientific output, Nuclear Science Symposium Conference Record, 2007. NSS '07. IEEE 2, Honolulu, HI USA, 2007 pp.1107-1109, ISBN 978-1-4244-0923-5.
- [3] Douglas Thain, Todd Tannenbaum, and Miron Livny, "Distributed Computing in Practice: The Condor Experience" Concurrency and Computation: Practice and Experience, Vol. 17, No. 2-4, pages 323-356, February-April, 2005
- [4] Journal of Physics: Conference Series Volume 119 Part 6, T Maeno 2008 J. Phys.: Conf. Ser. 119 062036 doi:10.1088/1742-6596/119/6/062036, PanDA: distributed production and distributed analysis system for ATLAS
- [5] <https://sdm.lbl.gov/twiki/bin/view/Software/BeStMan/WebHome>
- [6] Globus Toolkit Version 4: Software for Service-Oriented Systems. I. Foster. IFIP International Conference on Network and Parallel Computing, Springer-Verlag LNCS 3779, pp 2-13, 2006.
- [7] OSG Charter, <http://osg-docdb.opensciencegrid.org/cgi-bin/RetrieveFile?docid=25&extension=pdf>
- [8] National CI and the Campuses <http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=939>